Do You Think You Own Stock in the Power Company???

Michael Zyda, University of Southern California

In a previous "Games" column, I talked about the computational requirements of using OpenAI's Sora to potentially generate the complete Pixar 81-minute-long animated film Toy Story. *In this column, we discuss the power and water requirements for that potential computation.*

When I was a child, my father, Joseph Michael, used to admonish me when I left the lights on in a room with the bon mot, "do you think you own stock in the power company?" He was trying to get me to instantly turn off lights or whatever as I moved about our family's home. He also told me about his father, Sigmund, and how Sigmund turned off all lights, unplugged anything that was plugged in, lights, family radio, CPAP machine plugged in under the bed …, so that he could save whatever could be saved with such pecuniary minus any falls/missteps caused by the whole house being in the dark and air raid safe. I spent much of my childhood feeling the underside of the staircase's banister such that I could know precisely where I was and where to step despite the limitations of the darkness. I found myself going down the stairs with my eyes closed in the daytime so that I did not lose this capability. When I finally got to graduate school in St Louis, the Union Electric Company kindly included in each power bill the ability to purchase stock, which I did so that I could leave the lights on always and not trip over things in the dark. During my years in downtown Los Angeles, I had the joy of not having to leave the lights on as the helicopter beacons provided all that I needed for perambulation but no option for stock purchase at all.

So, my dear readers are probably wondering what the hell is this column about? Well, in July 2024, I came out with a column entitled "Can OpenAI's Sora Generate Pixar's *Toy Story*?," and I got buried in responses, I received two, that asked me how much power would it take to

Digital Object Identifier 10.1109/MC.2024.3487272
Date of current version: 26 February 2025

 0018-9162/25©2025IEEE

EDITOR **MICHAEL ZYDA**
University of Southern California;
zyda@mikezyda.com

perform that potential calculation? I thought I could just ignore my readership but every time I got into the swimming pool each morning it made me think about it as I bounced off each pool end. I had to try and provide an answer. And for some reason, I had it in my mind that I had to compare it to the amount of power that could be generated by a nuclear submarine, which left me with the problem of choosing which nuclear submarine as there are multiple classes. I had a retired Naval officer, former student of mine, now a patent attorney, let me know that the USS Nautilus generated 10 MW; the USS Nautilus first sailed in the year I was born, so maybe too dated a reference, and also provided the Los Angeles class submarines 150–165 MW and Seawolf class submarines 220 MW, the idea being that we could put our Nvidia server farm in a submarine underwater and reduce our cooling costs ...

How do we start figuring out the power requirements for the Pixar *Toy Story* computation?

Well, we need to first look back at Zyda[1] and the most important figure in that column, now our Figure 1.

In the article entitled "ChatGPT's resource demands are getting out of control" Andrew Tarantola[2] distilled a *Washington Post* article by Pranshu Verma and Shelly Tan entitled "A bottle of water per email: The hidden environmental costs of using AI chatbots"[3] with the goal of determining how much water and power would be used by asking ChatGPT to create a 100-word email. I am not sure why anyone would want ChatGPT to create a 100-word email—I get plenty of email already and spend mornings deleting most of it. I just know that I don't really need people getting the idea that they should be using ChatGPT and Nvidia farms to give us more spam-mail to read. Anyway, the *Washington Post* article Tarantola references determines that "using ChatGPT to write that 100-word email draws enough current to operate more than a dozen LED lightbulbs for an hour."[2] They also determine that water cooling of the Nvidia machine on which ChatGPT runs uses somewhere in the range of 235 mL to 1,408 mL of water for cooling. So, we are going to use these numbers and reference back to Zyda[1] to see how much power and water would be used should we be

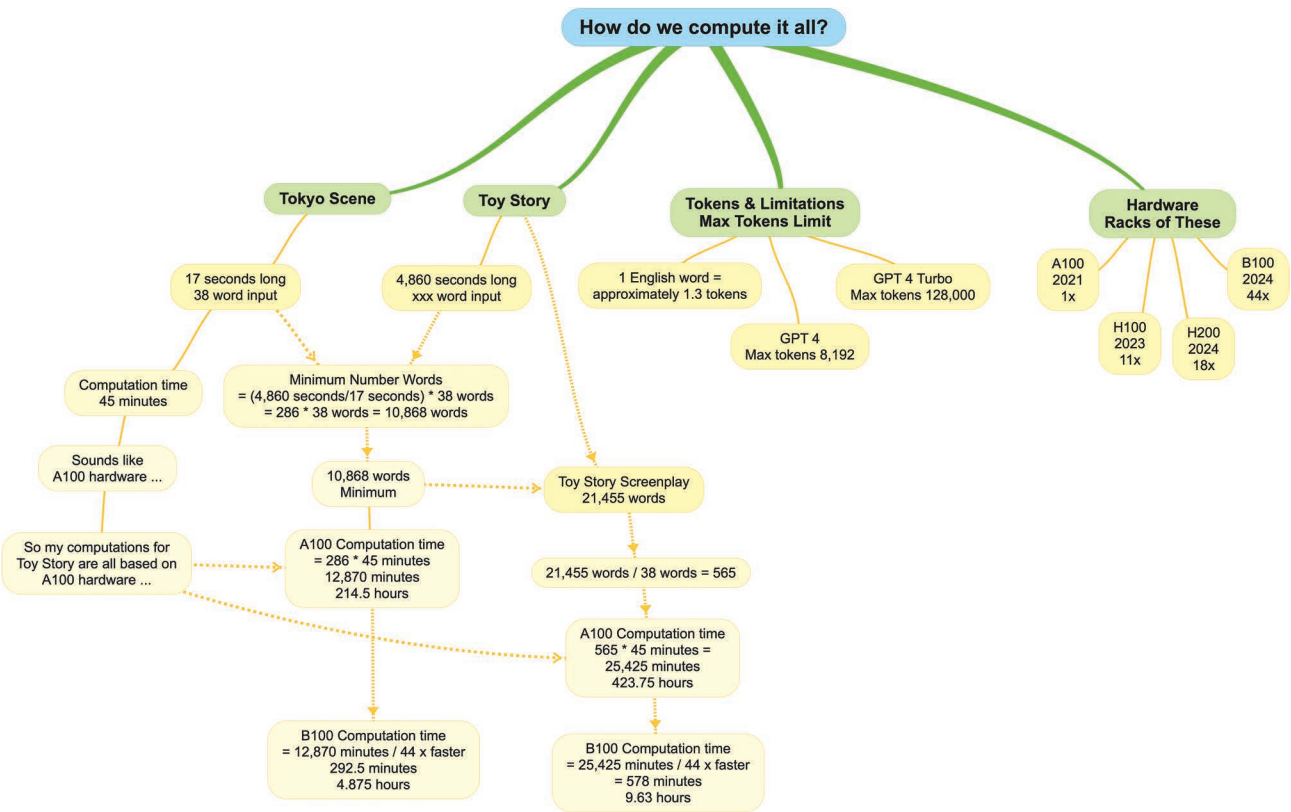

FIGURE 1. How do we compute it all?

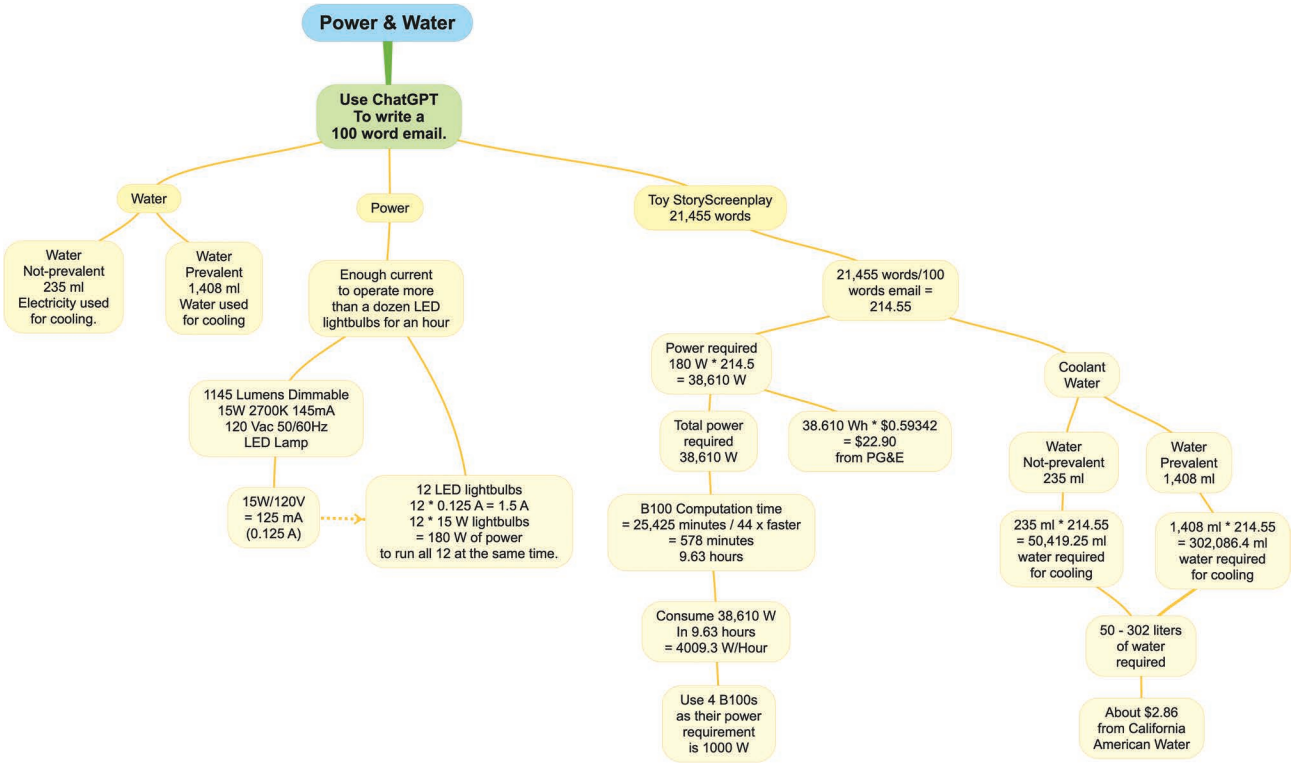

FIGURE 2. Power and water.

able to try and generate Pixar's *Toy Story* using OpenAI's Sora, which runs on top of ChatGPT.

COMMENTS?

If you have comments about this article, or topics or references I should have cited or you want to rant back to me on why what I say is nonsense, I want to hear. Every time we finish one of these columns, and it goes to print, what I'm going to do is get it up online and maybe point to it at my Facebook (mikezyda) and my LinkedIn (mikezyda) pages so that I can receive comments from you. Maybe we'll react to some of those comments in future columns or online to enlighten you in real time! This is the "Games" column. You have a wonderful day.

POWER AND WATER

So, in Zyda,[1] we used the *Toy Story* screenplay of 21,455 words as the number of words we needed to hand to Sora to have it generate the 81-min ersatz *Toy Story* of our technological swimming dreams. So we are going to take those 21,455 words and divide by 100-words and use that result, 214.55, to multiply against the 12 LED lightbulbs mentioned in Tarantola.[2] For our prototype lightbulb, we used the LED lightbulb closest to me in my basement office, a 15-W LED lightbulb, and multiplied by 12 to get 180 W of power to run all 12 lightbulbs. We then multiplied 180 W by the 214.55 result computed previously to get the total power required to generate the ersatz *Toy Story*, 38,610 W of required power. We then divided that power requirement by the 9.63 h of Sora computation time and find that we need something on the order of 4009.3 W per hour to be able to run this. We also know that the Nvidia B100 lists 1,000 W as its power requirement, so maybe we just run it across 4-B100s using our available 1,000-W power source. Clearly, we can run this from any of the aforementioned submarines. And the cost of that power from the Pacific Gas and Electric would be something on the order of US$22.90, which is way lower in cost than hiring all the engineers, artists, and designers that Pixar used to make the original 1989 film (Figure 2).

Using similar computations for the cooling water requirements, we find that 50 to 302 L of water are required, and that costs all of US$2.86 from the California American Water, if my conversions of gallons to liters were done properly. The real question that instantly came to my mind was why the Washington Post was using liters in an article whose primary audience is Americans who are prevented from using metric measures as a whole by oppressive political strictures.

I hope, dear readers, that this truncated column answers the posed questions of how much power and water are required to compute an ersatz *Toy Story* using OpenAI's Sora. There are some questions I did not answer but leave as an exercise to the erstwhile reader or two that demand such answers else they not be able to sleep. One that occurs to me is what does Nvidia do with the cooling water after it has cooled their fabulous hardware—do they just dump it into a river, stream, or ocean and thereby contribute to a more rapid onset of climate change? Don't they know that some of their key employees live in or come from states that recently have been drowned by the "change"? Perhaps we will have an answer by next column and see if we can stretch it out to a full memorable piece of technological literature.

ACKNOWLEDGMENTS

The author wishes to thank all of those readers who have gotten to the end of this bimonthly column without finding all of the deliberate and accidental errors. Those who wish to point out something earthshakingly seriously wrong with this column are not acknowledged—I will read your screeds and think about them while I am swimming.[4]

REFERENCES

1. M. Zyda, "Can OpenAI's Sora generate Pixar's toy story?," *Computer*, vol. 57, no. 7, pp. 136–140, Jul. 2024, doi: 10.1109/MC.2024.3393249.
2. A. Tarantola, "ChatGPT's resource demands are getting out of control," *Digital Trends*, Sep. 19, 2024. Accessed: Oct. 17, 2024. [Online]. Available: https://www.digitaltrends.com/computing/chatgpt-resource-demands/
3. P. Verma and S. Tan, "A bottle of water per email: The hidden environmental costs of using AI chatbots - AI bots generate a lot of heat, and keeping their computer servers running exacts a toll," *Washington Post*, Sep. 18 2024, Accessed: Oct. 17, 2024. [Online]. Available: https://www.washingtonpost.com/technology/2024/09/18/energy-ai-use-electricity-water-data-centers/
4. M. Zyda, *On Swimming*. (May 17, 2024). Accessed: Oct. 17, 2024. [Online Video]. Available: https://youtu.be/iP9NiJyYHs0

MICHAEL ZYDA is the founding director of the Computer Science Games Program and a professor emeritus of engineering practice in the Department of Computer Science, University of Southern California, Los Angeles, CA 90089 USA. Contact him at zyda@mikezyda.com.

Digital Object Identifier 10.1109/MC.2025.3533836